

RECEIVED
CENTRAL FAX CENTER

MAY 13 2008

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

First Named Applicant: New)	Art Unit: 2189
)	
Serial No.: 10/674,081)	Examiner: Peikari
)	
Filed: September 29, 2003)	HSJ920030174US1
)	
For: LOG-STRUCTURED FILE SYSTEM FOR DISK)	May 13, 2008
DRIVES WITH SHINGLED WRITING)	750 B STREET, Suite 3120
)	San Diego, CA 92101
)	

REPLY BRIEF

Commissioner of Patents and Trademarks

Dear Sir:

This Reply Brief is submitted in response to the Examiner's Answer dated April 29, 2008. The rebuttal in the Answer to Appellant's appeal brief begins on page 23 with an unusual statement, namely, that the examiner "will neither address nor concede aspersions cast upon the examiner with regard to fairness, impartiality..." Appellant points this out as an indication that, given that Appellant has never "cast aspersions" on the examiner's "fairness" or "impartiality", an inference is fairly raised that the examiner is persisting in the rejections potentially based on hypersensitivity to a perceived if non-existent slight.

On page 24 of the Answer the examiner maintains an unsupported definition of "check bytes" to be "data bytes", an overly broad construction that is necessary for the examiner to maintain his case but that fails to conform to MPEP §2111.01 (claim terms must be construed as the skilled artisan construes them). With more specificity, the evidence of record dispositively destroys the examiner's rejection by unambiguously

1189-14.R12

CASE NO.: HSI920030174US1

Serial No.: 10/674,081

May 13, 2008

Page 2

PATENT

Filed: September 29, 2003

demonstrating that the skilled artisan - the one relied on by the examiner - manifestly does not regard a check byte to be a data byte, contrary to the examiner's gambit. In paragraph 108 Asano teaches the following:

"...the ML-ISF-ECC [i.e., the Asano] process generates the ISF shared check bytes among N distinct 512 byte sectors *by summing stored check bytes and not the data bytes*" (emphasis mine).

If all check bytes were the same thing as data bytes as the examiner insists, then the above teaching *from the examiner's own reference* would be nonsensical. In effect, the examiner finds himself in a heated argument not so much with Appellant as with his own reference over what is and is not a "data byte". Appellant believes the evidence of record (the applied references as opposed to contradictory examiner argumentation) dispositively militates toward reversal.

On page 24 of the Answer the examiner continues to allege that "parity is probably the simplest example of an error correction code" and insists that "this is not an allegation, it is a fact." Appellant can only observe that (1) the examiner has yet to actually prove his "fact" with evidence of record, and (2) parity is not an error *correction* concept but an error *detection* concept (see, e.g., orafac.com, satech.com, notepage.net, and micro2000.co.uk, all of which define "parity" to be related to error *detection*). Detecting an error and correcting it are two different things. That's a fact.

On page 25 of the Answer the examiner insists that paragraphs 107-109 of Asano teach successive partial writes and an ECC block being retained. Appellant first observes that the actual claim language requires that a cumulative ECC parity state between successive partial writes of an ECC block is retained.

1189-14.NPL

CASE NO.: HSJ920030174US1

Serial No.: 10/674,081

May 13, 2008

Page 3

PATENT

Filed: September 29, 2003

Appellant will next play unfairly. Specifically, Appellant will copy in paragraphs 107-109 of Asano and let them speak for themselves:

"[0107] In summary, the present invention provides a solution to the problem of using a physical long block error correction code (ECC) encoding and decoding scheme by providing a low check byte overhead on the hard disk drive (HDD), for data protection against mixed mode errors, while maintaining a 512 byte sector logical unit compatible with present operating systems such as, for example, Microsoft Windows.RTM. and Linux.RTM., without necessarily incurring the read-modify-write (RMW) process performance penalty. The solution is most effective for those storage applications that typically involve large block transfers, such as, for example, audio-visual (AV) applications, such as streaming AV data. The solution protects most of the data on the disk against large burst errors while maintaining on-the-fly (OTF) ECC protection against random errors for all of the data on the disk, without incurring the RMW process performance penalty.

"[0108] The ML-ISF-ECC process is a solution that permits the use of a long block, multiple-sector, multiple level ECC, at a low average check byte overhead, without requiring a change in the operating system standard and which does not require RMW. A fixed number, N, of 512 byte sectors are integrated in a multiple level physical multi-sector block ECC by a multitude of encoders whose check bytes are stored and summed appropriately such as to generate shared multiple level check bytes, protected by the OTF ECC, that can be employed for correction in the integrated long block of sectors, as required by the individual sector failures. Unlike the integrated interleaving technique disclosed in U.S. Pat. No. 5,946,328, where simultaneous byte-by-byte summation of individual short codewords is required for the generation of shared check bytes, the ML-ISF-ECC process generates the ISF shared check bytes among N distinct 512 byte sectors by summing stored check bytes and not the data bytes. The check bytes obtained from encoding each individual sector by a multitude of encoders are stored and summed appropriately after all N integrated sectors have been separately encoded. Furthermore, with the ML-ISF-ECC process there is no limit on the number of levels of correction.

"[0109] When the number of sectors required to be written on the hard disk by the host operating system, which provides the initial and last logical block address (LBA), is not an integer multiple of N logical 512 byte sector units, the multiple level encoding operation may be partially disabled leaving only those levels of ECC, compatible with fragmentation, enabled. The performance degradation due to RMW is thus avoided at the expense of reducing the ML-ISF-ECC from a multiple level ECC to a first level OTF ECC for those long blocks that are fragmented. Specifically, for AV applications this is an excellent tradeoff as most data

1195-14.RPI.

CASE NO.: HSI920030174US1
Serial No.: 10/674,081
May 13, 2008
Page 4

PATENT
Filed: September 29, 2003

transfers involve very large data blocks and thus the fraction of the data not protected by the ML-ISF-ECC process will be very small."

Thus, at its most relevant Asano teaches in these portions that a fixed number of sectors are integrated in a multiple level physical multi-sector block ECC by a multitude of encoders whose check bytes are stored and summed appropriately such as to generate shared multiple level check bytes. Compare this to the relevant language of Claim 1: "a cumulative ECC parity state between successive partial writes of an ECC block is retained". There is no mention of "successive partial writes" in the relied-upon parts of Asano or any cognizable synonym, much less in a way that would meet this claim limitation. Appellant respectfully suggests that reversal is appropriate.

With respect to Appellant's observation that no evidence exists that Rosenblum's log-structure file system would indeed result in faster operation in an application Rosenblum never contemplates, the examiner uses the expedient of assuming that since Rosenblum is designed in general for disks, it must always and everywhere suggest itself in every disk application that has been or ever shall be. The problem here is almost the polar opposite of the situation in KSR Int'l Co. v. Teleflex Inc., 127 S.Ct. 1727 (2007), in which "predictable variations" were at issue. In contrast, Rosenblum has not, on the evidence, been recognized as being effective or even useful as a vehicle that uses the ECC system of Asano. The only place that any such relevant recognition exists on the present record is in the present specification on page 8 that such a log structure might find use in the context of Claim 1:

Appellant is not setting "forth an unusual standard for obviousness" (Office Action, various places including top of page 29). Appellant is merely observing that lifting motivations from a reference that apply

1189-14.RPI,

CASE NO.: HSI920030174US1

Serial No.: 10/674,081

May 13, 2008

Page 5


PATENT

Filed: September 29, 2003

to one field and using them as rationales for combining the reference with another reference are insufficient unless the evidence demonstrates that such a *combination* would have been a "predictable variation", KSR. That is, the proffered motivation to combine must find relevance in the end combination, not merely in the particular field of an individual reference; otherwise, since every patent extols its virtues in a vacuum, every patent would be combinable with every other patent. That would be an unusual standard.

Concluding with something truly unusual, a patent examiner *and two senior members of the examining corps* have placed on the written record a declaration that analysis under MPEP §2173.02 "is useless unless appellant can provide an unambiguous explanation of how a "state" can be made "cumulative"" (Answer, page 31, lines 13-15). Appellant herewith defends the MPEP. The analysis of Section 2173.02 is not "useless" pending an examiner-demanded condition precedent; it is not even discretionary, but mandatory. The declaration of the uselessness of MPEP guidance relative to the indefiniteness rejection would seem to justify reversal on that ground alone, indelibly proving, as it does, that the conferees have willfully ignored the requirements in the MPEP for making their case.

Respectfully submitted,



John L. Rogitz
Registration No. 33,549
Attorney of Record
750 B Street, Suite 3120
San Diego, CA 92101
Telephone: (619) 338-8075

JLR:jg

1100-14 JLR